



**BERKELEY LAB**  
Bringing Science Solutions to the World



U.S. DEPARTMENT OF  
**ENERGY**

## SENSEI: Cross-Platform View of In Situ Analytics

E. Wes Bethel, Junmin Gu, Burlen Loring, Dmitriy Morozov, Gunther H. Weber, John Wu (LBNL). Nicola Ferrier, Joseph Insley, [Silvio Rizzi](#), [Sergei Shudler](#) (ANL). Dave Pugmire, James Kress, Matthew Wolf (ORNL). Earl Duque, Brad Whitlock (Intelligent Light). Utkarsh Ayachit, David Thompson, Andrew Bauer, Patrick O'Leary (Kitware)



# Acknowledgment

This work is supported by the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. Department of Energy, Office of Advanced Scientific Computing Research, under Contract No. DE-AC02-05CH11231, through the grant “**Scalable Analysis Methods and *In Situ* Infrastructure for Extreme Scale Knowledge Discovery,**” program managers Dr. Lucy Nowell and Dr. Laura Biven.



SENSEI: Scalable Analysis Methods  
and *In Situ* Infrastructure for Extreme  
Scale Knowledge Discovery



- Today we will cover a subset of a half day tutorial presented at SC18 by the SENSEI team
- Download SC18 slides and virtual machine image:  
<https://sensei-insitu.org/tutorials/sc18.html>



The screenshot shows the SC18 website interface. At the top is a blue navigation bar with the SC18 logo and menu items: PROGRAM, EXHIBITS, EXPERIENCE, and SUBMIT. Below the navigation bar, a breadcrumb trail shows 'Presentation'. The main heading is 'Presentation' in a large red font. A horizontal menu of links includes FULL PROGRAM, PRESENTERS, ORGANIZATIONS, SEARCH PROGRAM, FLAGGED, HAPPENING NOW, MAPS, and NOTIFICATIONS. The featured event is 'SENSEI Cross-Platform View of In Situ Analytics'. Below the title, the presenters are listed: E. Wes Bethel, David Thompson, Burlen Loring, Silvio Rizzi, Brad Whitlock, Matthew Wolf, and Patrick O'Leary. The event type is 'Tutorial'. The registration category is 'TUT', shown in a red box. The tags are 'Data Analytics', 'Data Management', 'Productivity', and 'Visualization'. The time is 'Sunday, November 11th, 1:30pm - 5pm', with small icons for social media or sharing.

## Relevant links

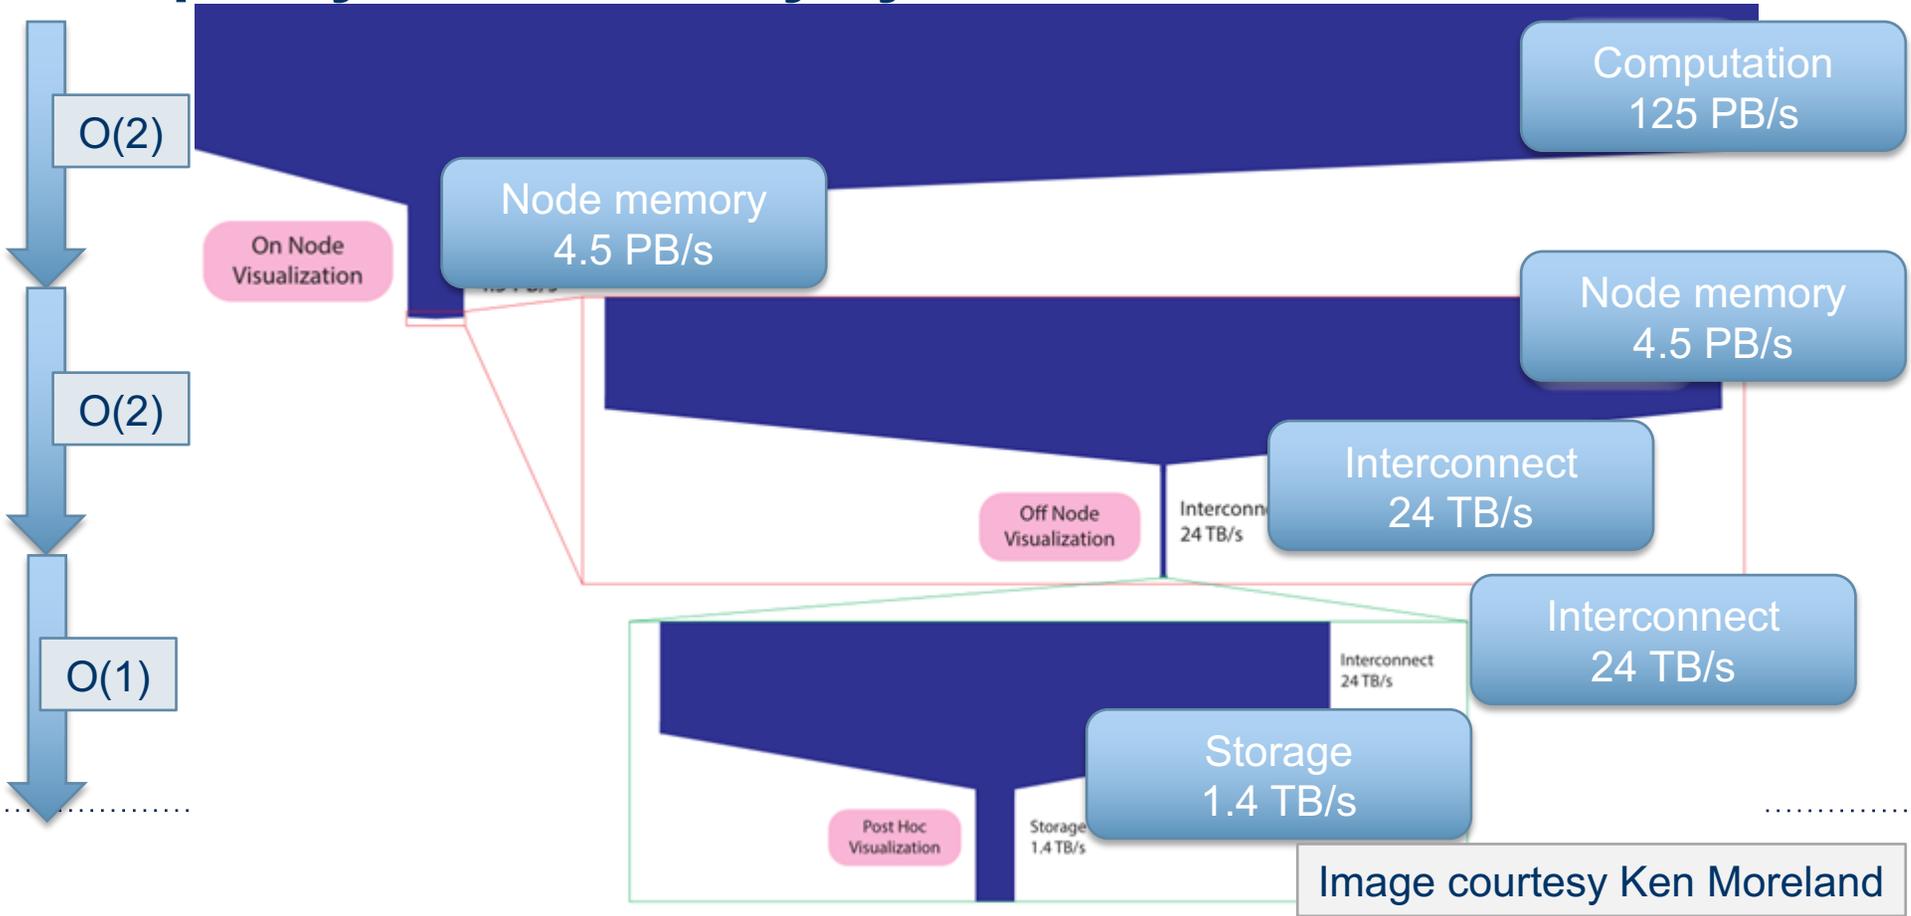
- Main page – <http://www.sensei-insitu.org/>
- Software repository – <https://gitlab.kitware.com/sensei/sensei>
- ADIOS – <https://www.olcf.ornl.gov/center-projects/adios/>
- VisIt/Libsim – <https://www.visitusers.org/index.php?title=Category:Libsim>
- ParaView Catalyst – <http://www.paraview.org/in-situ/>
- SENSEI in situ tutorial at CSCS

<https://www.youtube.com/watch?v=nA22JqzhjqQ&list=PL1tk5lGm7zvRSS-M2bvW3JCt93gpgGHjM>

---



# Five orders of magnitude between compute and I/O capacity on Titan Cray system at ORNL



## What is *in situ* data analysis and visualization?

- **Post processing**: save to disk, then later, a separate analysis/vis program reads that data and operates on it.
  - **In situ processing**: process data as it produced without writing to and reading from storage. Processed “in place”.
    - Many flavors/terms: tightly coupled, loosely coupled, in transit, co-processing, etc.
    - Practical view: anything processed but not written to persistent storage is *in situ*
-

## Generic processing sequence

1. initialize sim
  2. do
  3.   compute new state
  4.   if do\_io write plot file
  5. while !done
  6. finalize sim
-

## Generic processing sequence w/ in situ

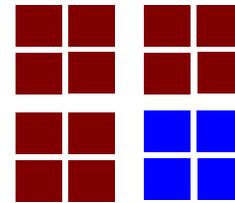
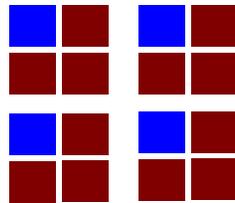
1. initialize sim
2. **if do\_insitu initialize in situ**
3. do
4.   compute new state
5.   if do\_io write plot file
6.   **if do\_insitu execute in situ**
7. while !done
8. **if do\_insitu finalize insitu**
9. finalize sim

execute is where things get interesting

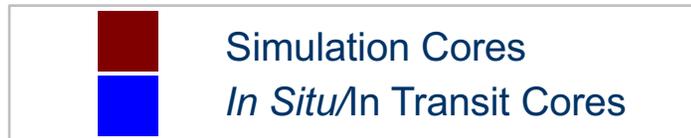
- shared address space zero copy data transfers to shared or unique compute resources
  - staging transfer sends data to a de-coupled parallel job, potentially asynchronous, potentially different jobs size
-

# In situ vs In transit

In situ – no data movement:  
Simulation and *in situ* methods share memory



In transit – data is moved:  
Simulation and *in situ* methods do not share memory



## What is the cost of *in situ* processing?

Concern: simulations want to use all available resources, so having an understanding of *in situ* resource utilization is useful.

In other words: In situ infrastructure must play nicely with simulation

Full details in SC16 paper:

Ayachit, Bauer, Duque, Eisenhauer, Ferrier, Gu, Jansen, Loring, Lukic, Menon, Morozov, O'Leary, Ranjan, Rasquin, Stone, Vishwanath, Weber, Whitlock, Wolf, Wu, and Bethel, "*Performance Analysis, Design Considerations, and Applications of Extreme-scale In Situ Infrastructures*". In Proceedings of SC16, November 2016.

---



**BERKELEY LAB**  
Bringing Science Solutions to the World



U.S. DEPARTMENT OF  
**ENERGY**

# SENSEI System Overview



## *In situ* infrastructures

Relatively new. Until recently, **ad hoc**, **proof-of-concept prototypes**. However, several **production quality *in situ* infrastructures** have emerged

**ADIOS** provides tools for *in situ* **I/O** , **data movement** and **analysis**

- ADIOS allows simulations to adopt *in situ* techniques by **leveraging** their **advanced I/O infrastructures** that enable co-analysis pipelines **rather than changing the simulator**.
- The non-intrusive integration **provides resilience** to third party library bugs and possible jitter in the simulation.

**ParaView** and **VisIt** both provide tools for *in situ* **analysis** and **visualization**

- Can be **tightly** or **loosely** linked to a simulation, allowing the simulation to **share data** with Catalyst for analysis and visualization.
- Catalyst, Libsim, and ADIOS enable the **opposite flow of information**, sending data from the client to the simulation, enabling the possibility of *in situ* and/or **monitoring/simulation steering**.

**Ascent** an emerging *in situ* framework with an elegant data model, taking advantage of emerging **VTK-m** many core analysis and rendering capabilities

---

## Can WE....

Enable use of any in situ framework?

Enable use of any analysis library/tool, even those not designed for in situ?

Develop analysis routines that are portable between codes?

Make it easy to use?

---

## The *current* problem set



SENSEI seamlessly & efficiently enables in situ data processing with a diverse set of tools & libraries

---

# Our approach

## Data model

- The lingua franca allowing an analyses to access simulation data consistently across a variety of simulations

## Data adaptor

- Convert simulation data to/from the data model
- API for accessing the simulation data from the backend

## Analysis adaptor

- Present the back-end data consumer to the simulation
- API for pushing data through the system from the sim

## Library

- Providing off the shelf access to a diverse set of back-ends. eg Libsim, Catalyst, and ADIOS capabilities
- 



## Write once run everywhere

The **SENSEI API** enables connection of simulation data sources to visualization and analysis back ends

- From the perspective of the simulation, the back ends (analysis/vis codes) are interchangeable

The **SENSEI data model** enables viz & analysis codes to access data through a unified API.

- From the perspective of the analysis/visualization code, data sources (simulations) are interchangeable
-

# SENSEI Architecture

“write once  
A simulation  
end through  
the back

SENSEI’s analysis  
adaptors provide the  
API for simulations to  
drive analysis and vis

Simulation



Catalyst  
adaptor



Lisim  
adaptor



ADIOS  
adaptor



Python  
adaptor



SENSEI’s data adaptor  
API and data  
expose simulation  
structures to the  
back-end

Bridge/instrumentation code  
is added to call SENSEI  
Analysis. Typically: Initialize,  
Execute, Finalize

```
1. initialize sim
2. if do_insitu bridge::initialize
3. do
4.   compute new state
5.   if do_io write plot file
6.   if do_insitu bridge::execute
7. while !done
8. if do_insitu bridge::finalize
9. finalize sim
```



**BERKELEY LAB**  
Bringing Science Solutions to the World



U.S. DEPARTMENT OF  
**ENERGY**

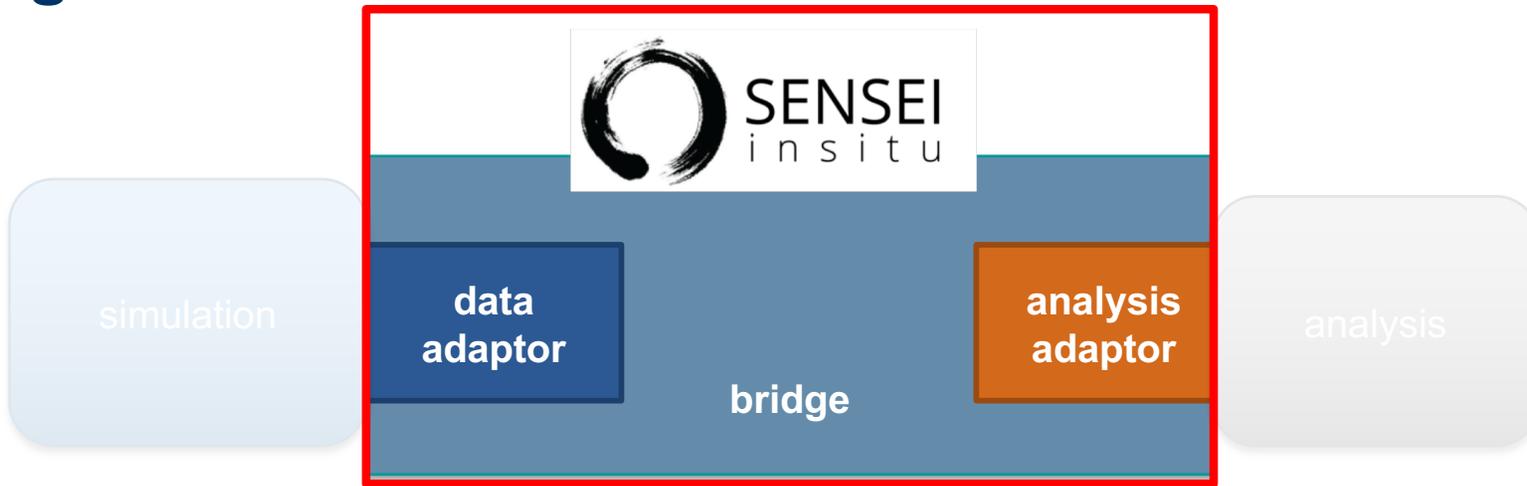
## SENSEI API's



**Intelligent Light**



## Bridge API

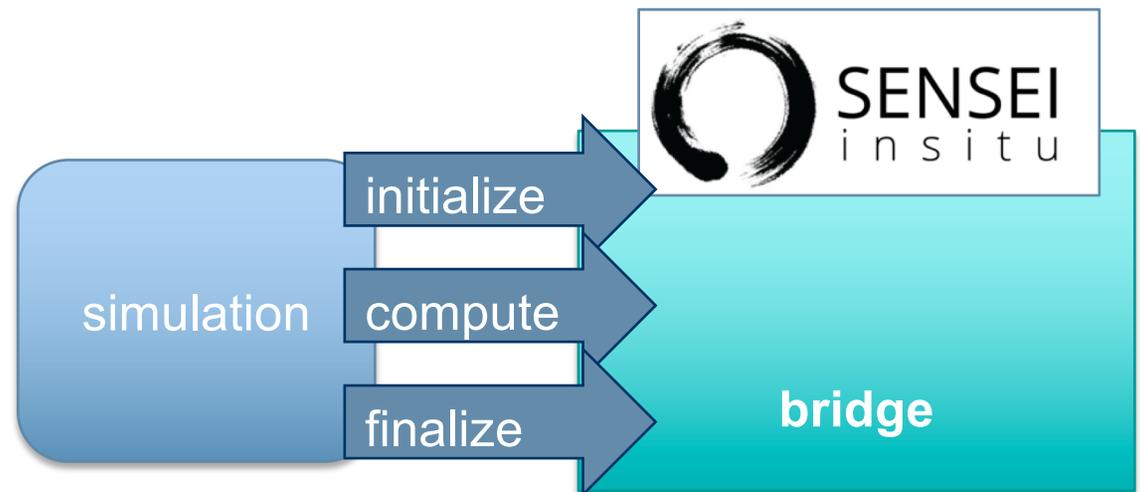


- Is part of the simulation code
  - Is where you create, initialize, and manage your data and analysis adaptors
  - Is where you execute the analyses adaptors as needed
  - Typically consists of 3 functions: Initialize, Compute and Finalize
-

# Implementing the bridge to SENSEI

Typically 3 calls:

- Initialize()
    - Set the DataAdaptor
    - Initialize DataTimeStep
    - Specify what analysis will be done. For the Oscillator we use the ConfigurableAnalysis class.
  - Compute()
    - For the Oscillator we do this with two calls: set\_data() / set\_particles() and analyze(), so that SENSEI may be disabled in benchmarks
  - Finalize()
- 



## Simulation loop with bridge code

1. initialize sim
  2. **if do\_insitu bridge::initialize**
  3. do
  4.   compute new state
  5.   if do\_io write plot file
  6.   **if do\_insitu bridge::execute**
  7. while !done
  8. **if do\_insitu bridge::finalize**
  9. finalize sim
-

# Run time configuration

## Adaptors

- SENSEI Configurable analysis. Parses XML and creates and configures one of the other analysis adaptors interfacing to the back-ends (Libsim, Catalyst, ADIOS, custom, etc).
- Direct integration

## Back-ends

- May expose control API via their SENSEI adaptor. In the Configurable analysis adaptor these are exposed via XML attributes.
  - May be scriptable via their own Python bindings adding another layer of control.
  - May be configured via "state" or "session" files.
  - Special purpose
-

# ConfigurableAnalysisAdaptor

- A meta analysis. A manager. It configures and invokes one or more of the other analysis adaptors
  - XML specifies analyses and their run time options
  - Supports ADIOS, Catalyst, Libsim, VTK I/O, and other data consumers
  - In in transit use cases one XML configures the transport a second configures the analysis/backend
-

# ConfigurableAnalysis XML

```
<sensei>
  <!-- Custom Analyses -->
  <analysis type="histogram" mesh="bodies" array="v" association="point"
    bins="10" enabled="0" />

  <!-- VTK XMLP I/O -->
  <analysis type="PosthocIO" mode="paraview" output_dir="." enabled="0">
    <mesh name="bodies">
      <point_arrays> ids, m, v, f </point_arrays>
    </mesh>
  </analysis>

  <!-- CATALYST -->
  <analysis type="catalyst" pipeline="pythonscript"
    filename="../sensei/miniapps/newton/newton_catalyst.py" enabled="1" />

  <!-- LIBSIM -->
  <analysis type="libsim" plots="Pseudocolor" plotvars="ids"
    image-filename="newton_%ts" image-width="800" image-height="800"
    slice-project="1" image-format="png" enabled="0"/>
</sensei>
```

---



**BERKELEY LAB**  
Bringing Science Solutions to the World



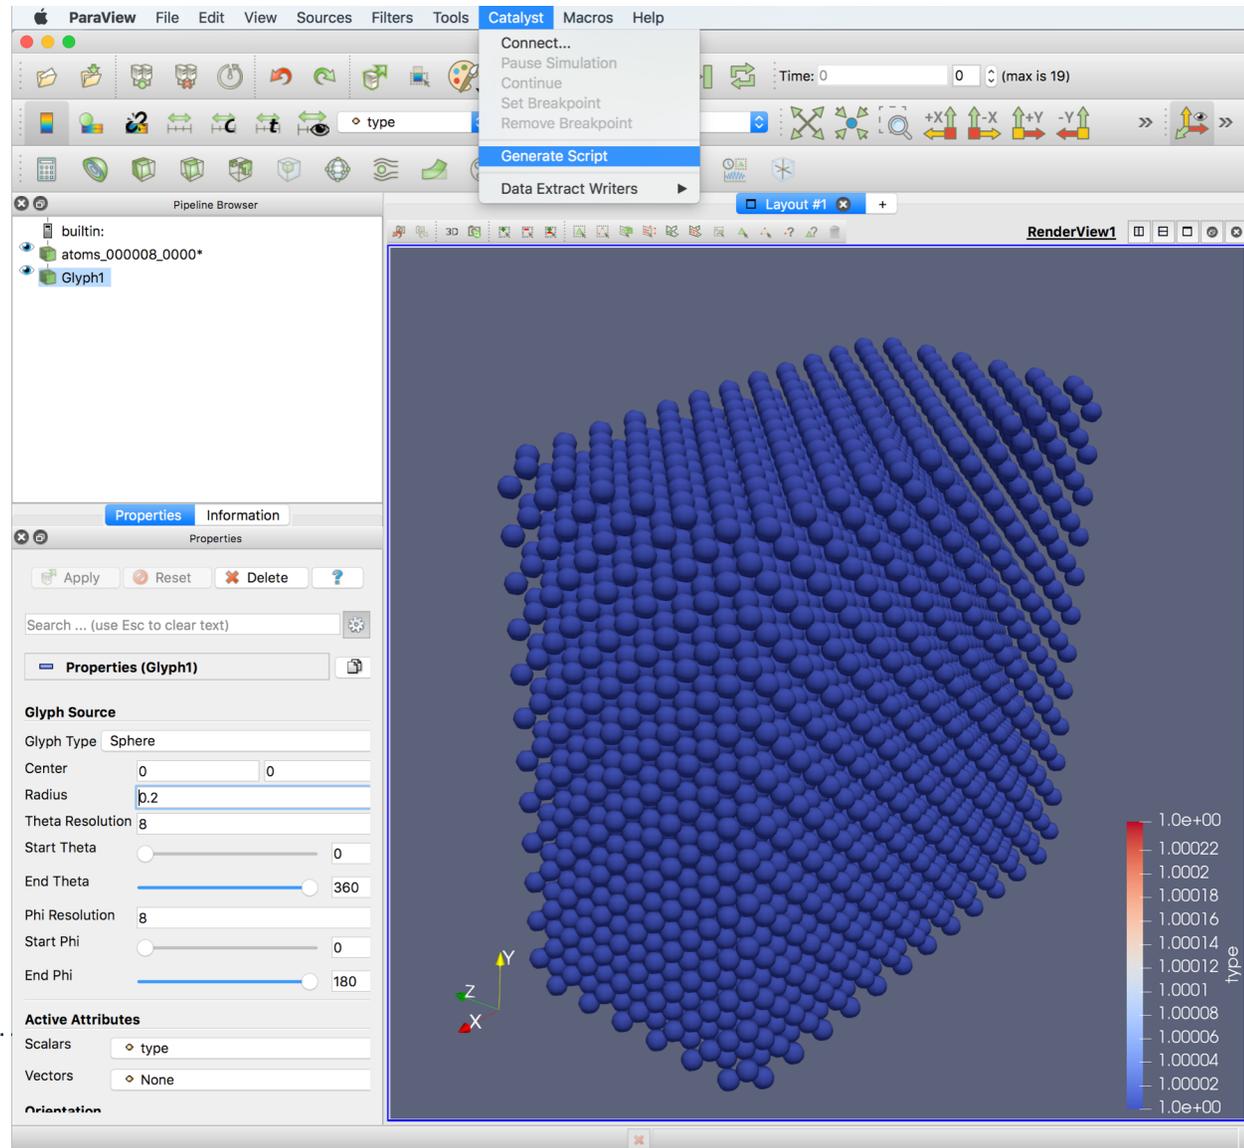
U.S. DEPARTMENT OF  
**ENERGY**

## Generating an in situ pipeline in ParaView



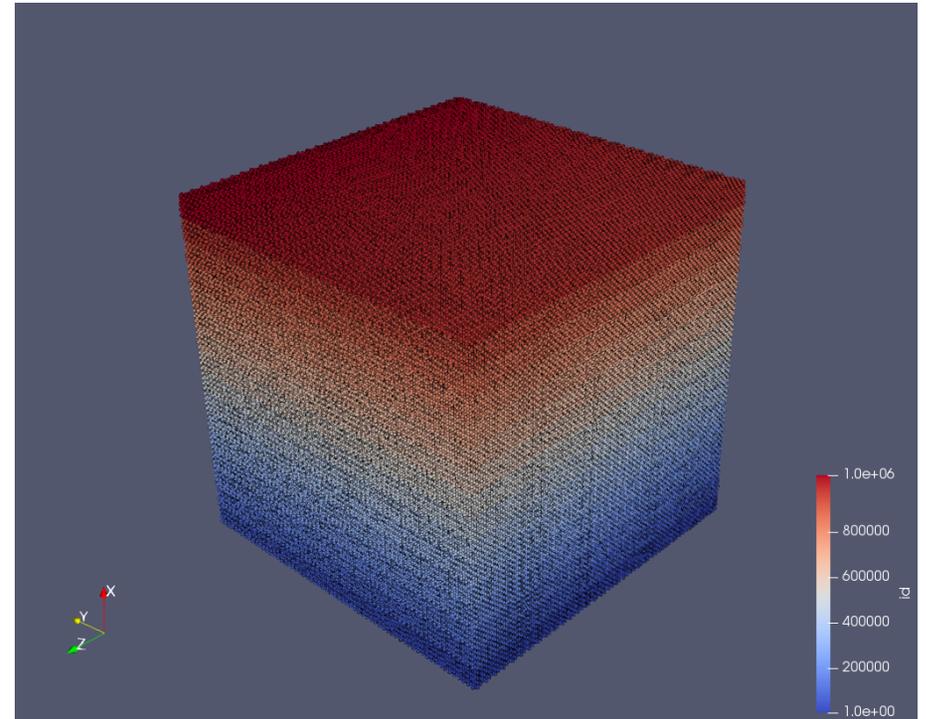
# Instructions for ParaView v5.5.2

- Load a representative dataset in ParaView
- Define your visualization pipeline
- Export Catalyst Python script



## Catalyst example

- Configure XML file
- Run instrumented simulation
- Result: one .png image per simulation timestep



```
<sensei>  
  <!-- Available with ENABLE_CATALYST -->  
  <analysis type="catalyst" pipeline="pythonscript"  
    filename="gaussianptsbyid.py" enabled="1" />  
</sensei>
```



**BERKELEY LAB**  
Bringing Science Solutions to the World



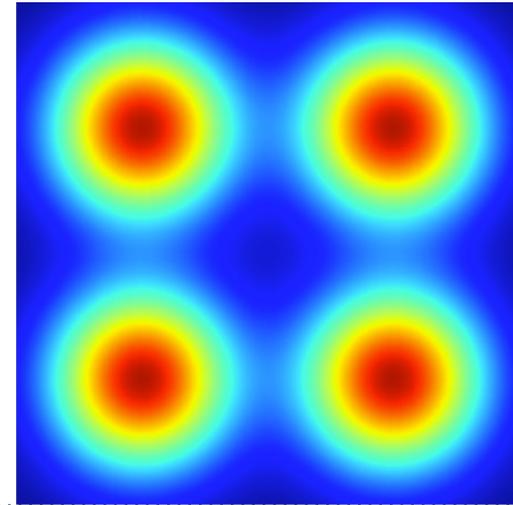
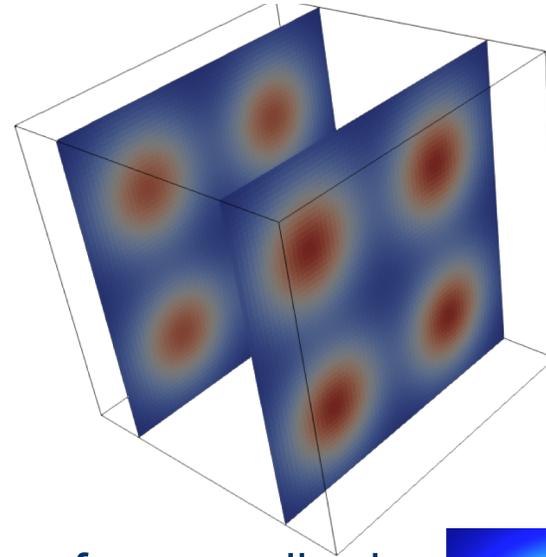
U.S. DEPARTMENT OF  
**ENERGY**

## Demos

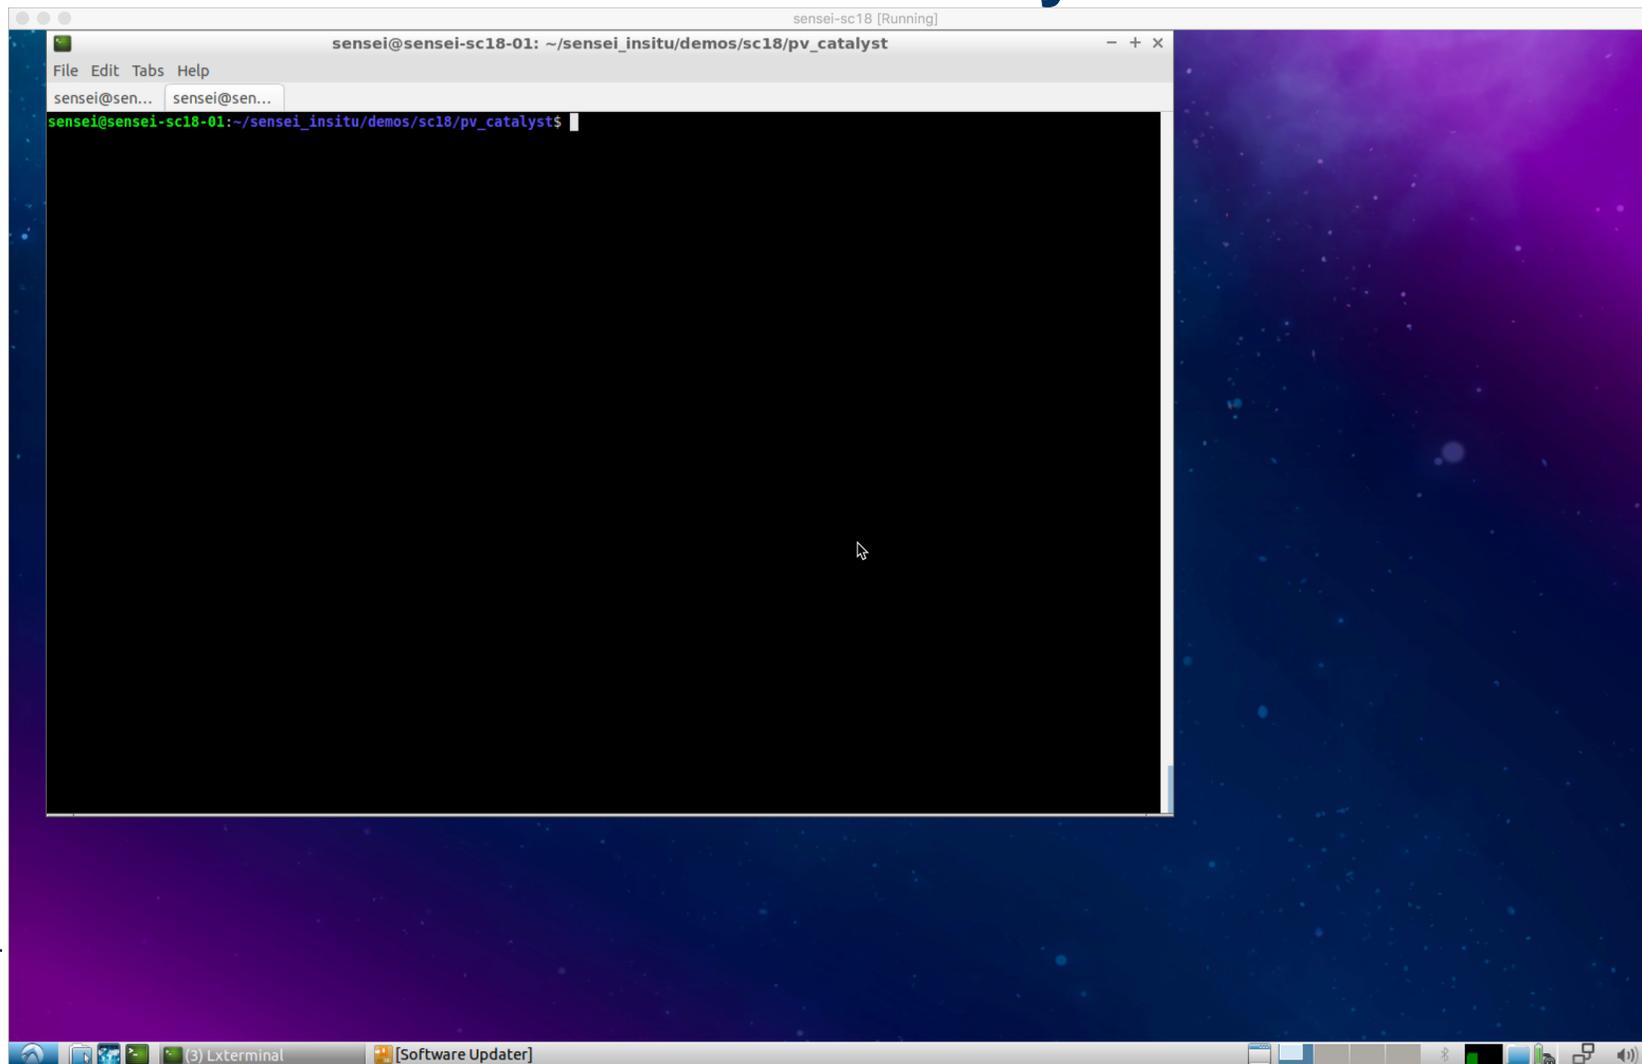


## Oscillator miniapp overview

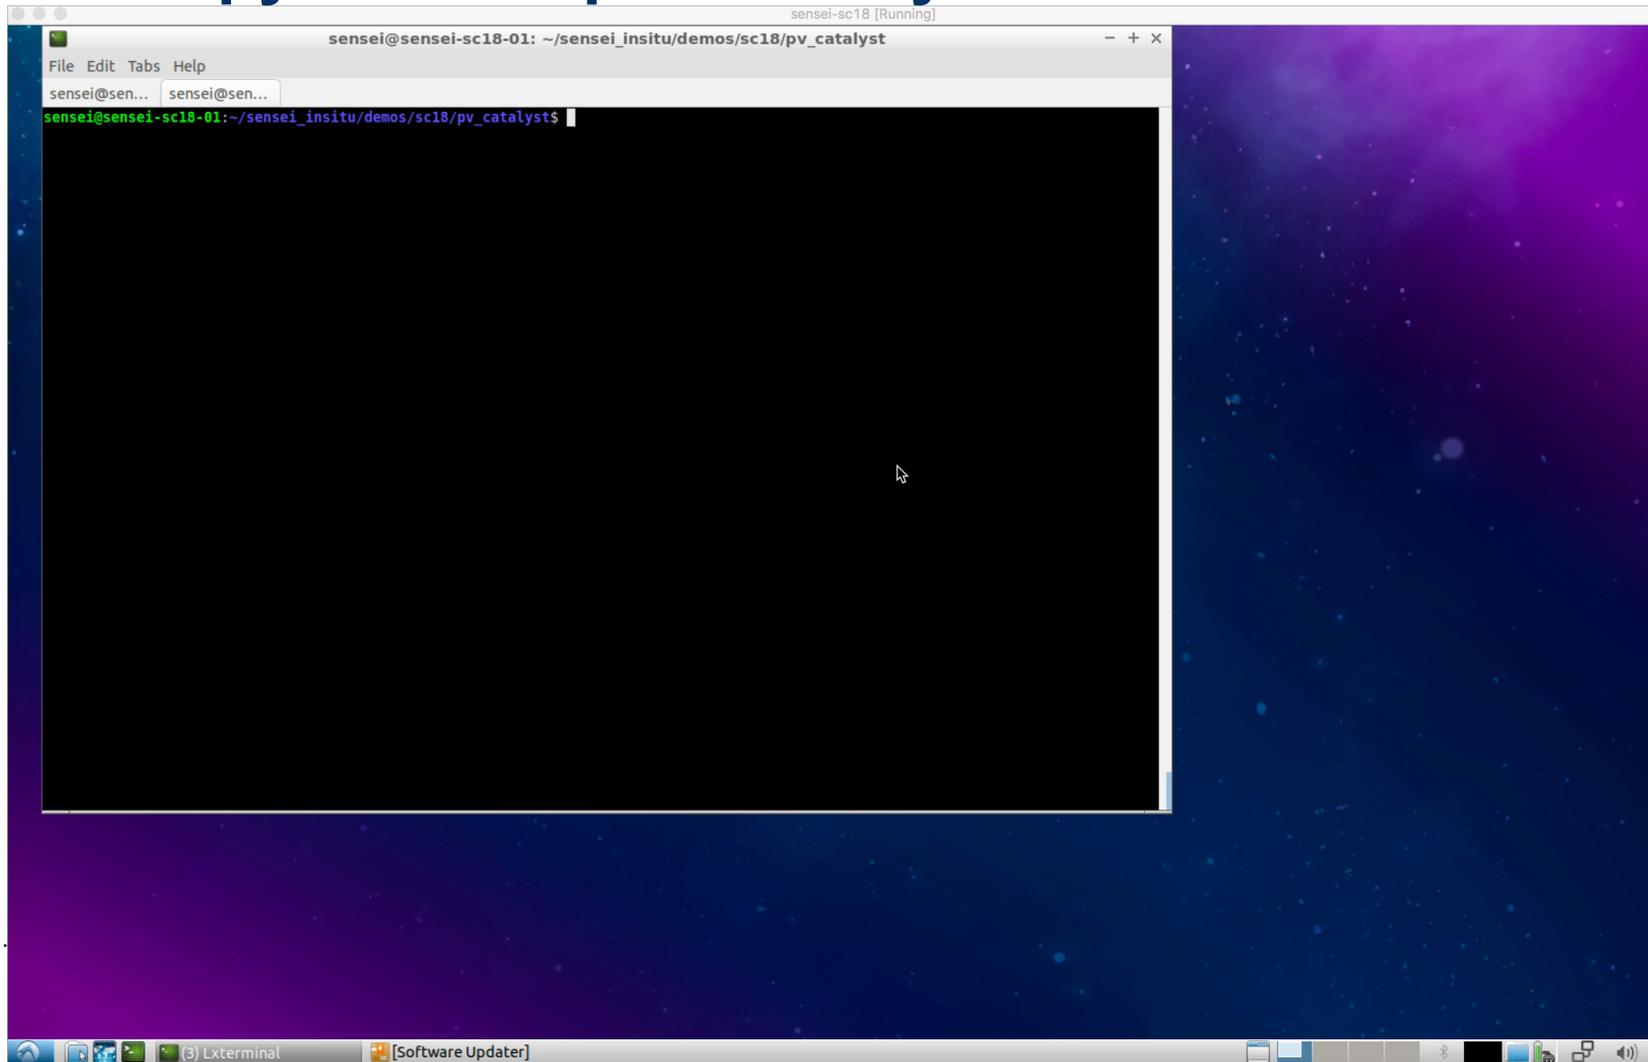
- MPI based C++ code that simulates a collection of periodic, damped, or decaying oscillators over a Cartesian grid.
- Unstructured grid also supported
- Each oscillator is convolved with a Gaussian of a prescribed width
- Can randomly place particles and advect them using an analytical velocity field
- Executable inputs are oscillator parameters, time resolution, length of the simulation, grid dimensions, grid partitioning, and number of random particles to generate



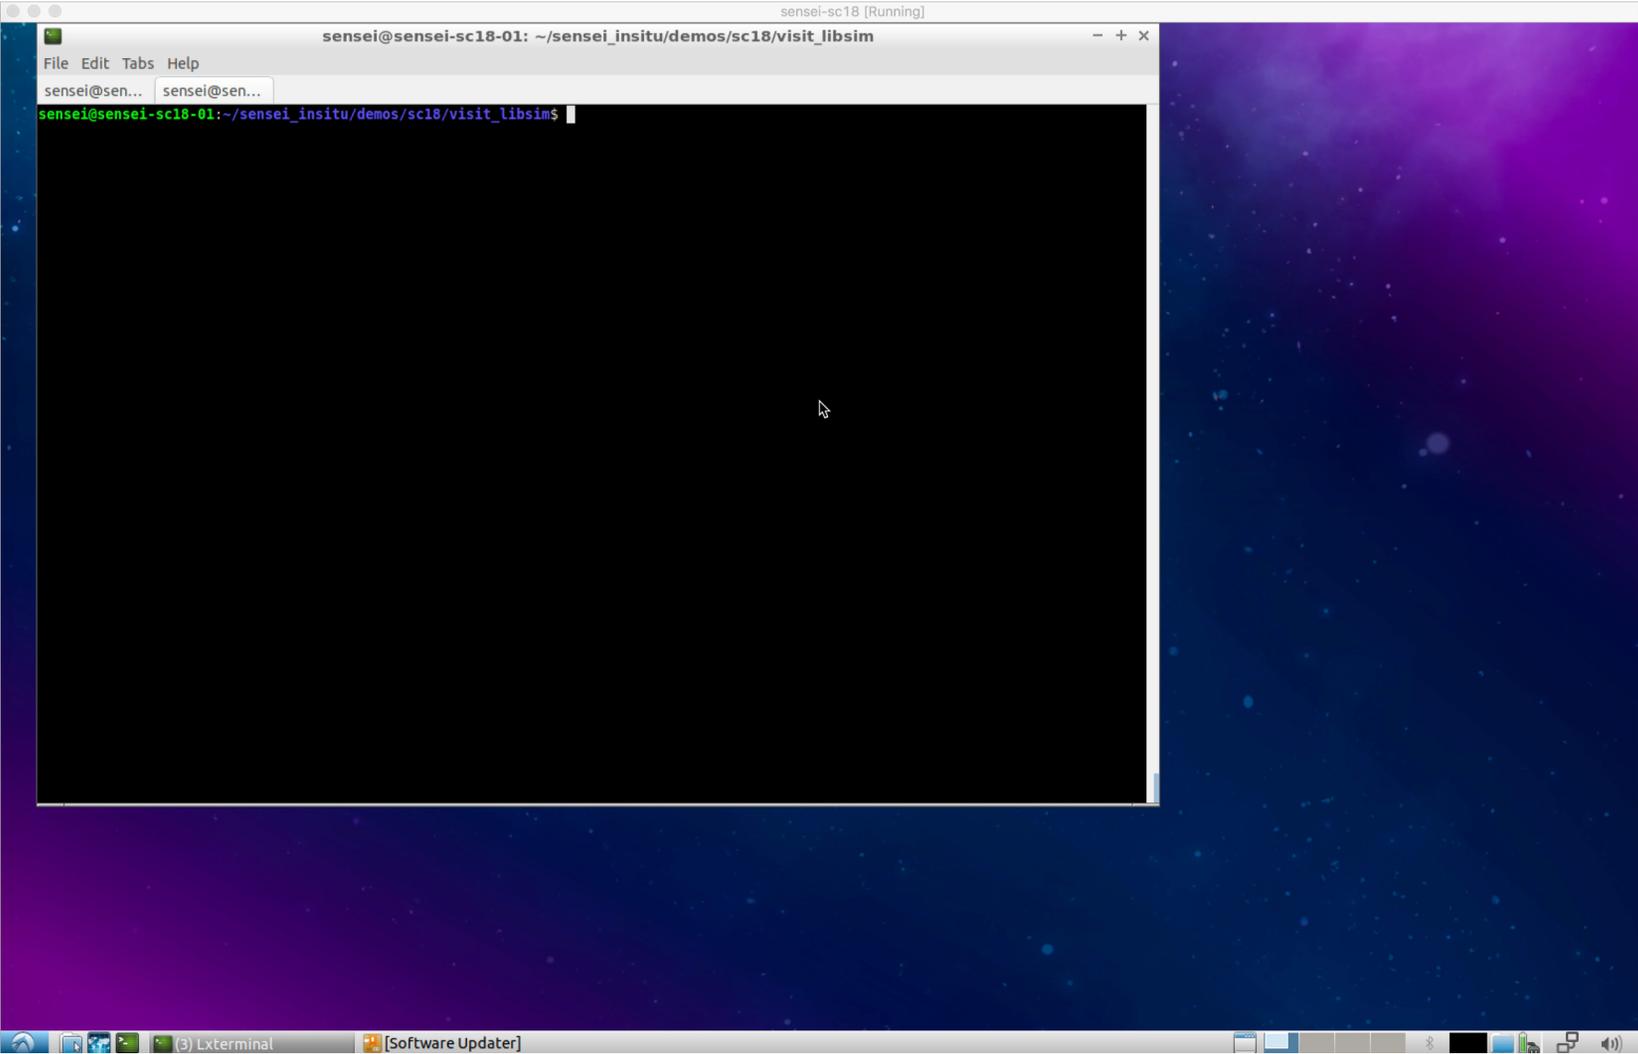
# Demo 1: create slices with Catalyst in situ



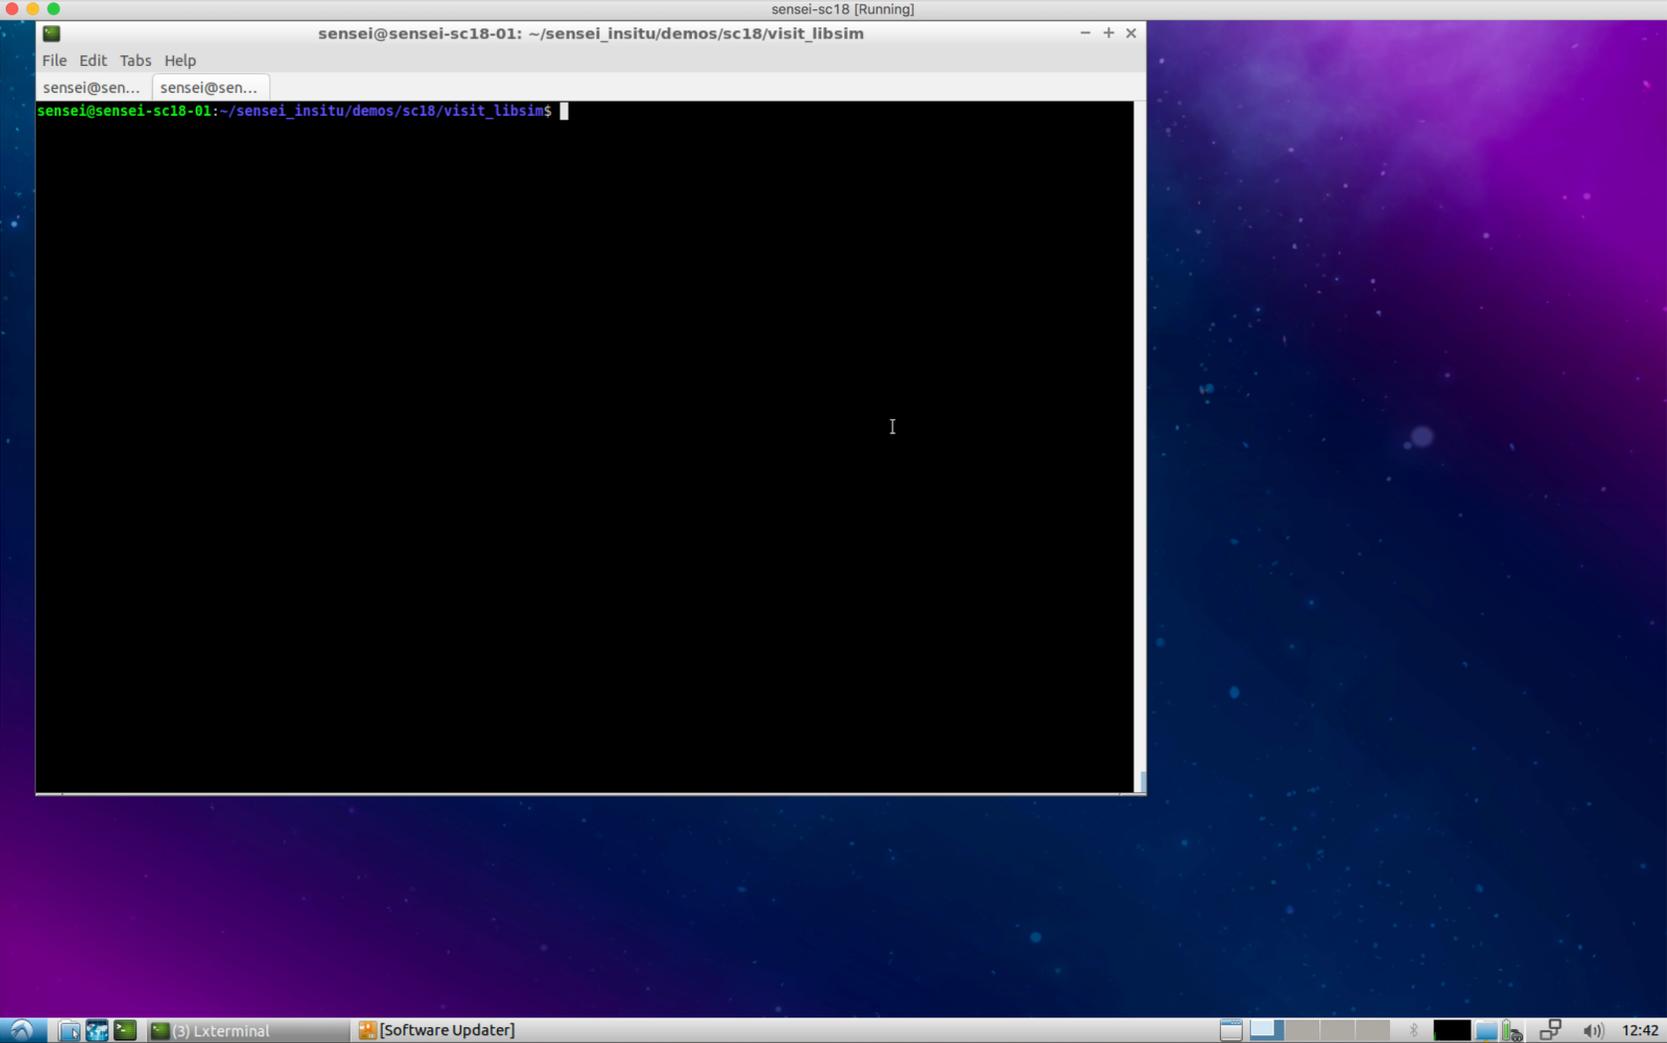
## Demo 2: python script in Catalyst in situ



# Demo 3: unstructured mesh with Libsim in situ



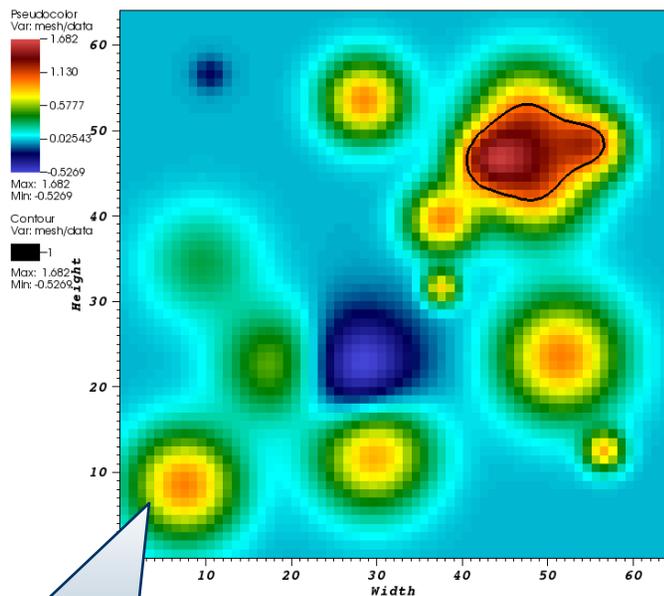
# Demo 4: exploring Libsim extracts with Visit in situ



# More demos in the SC18 tutorial

- ADIOS in transit demo
- Python mini app instrumented with SENSEI
- Catalyst, VTK-m, Haar wavelet, and Cinema databases

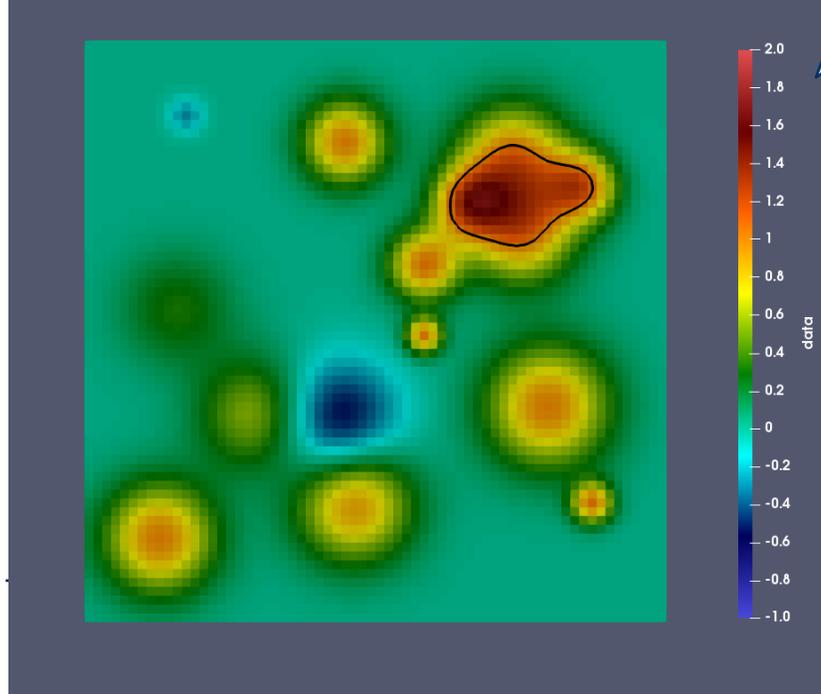
DB: batch.sim2  
Cycle: 79 Time:3.95



rendered with libsim

user: sengel  
Tue Oct 30 15:50:53 2018

$t = 3.95$



rendered with catalyst



**BERKELEY LAB**  
Bringing Science Solutions to the World



U.S. DEPARTMENT OF  
**ENERGY**

# Instrumenting LAMMPS with SENSEI

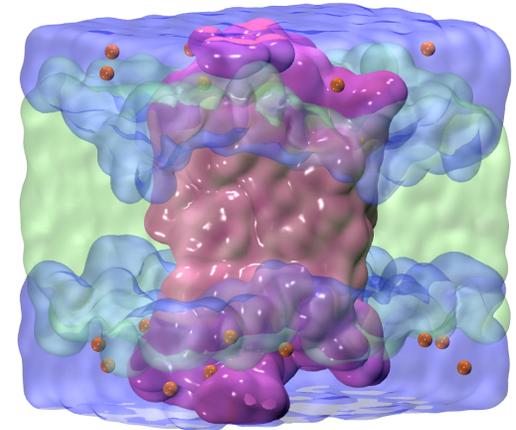


# LAMMPS

- Large-scale Atomic/Molecular Massively Parallel Simulator
- Classical molecular dynamics code
- Runs on single processors or in parallel using message-passing techniques and a spatial-decomposition of the simulation domain
- Accelerated performance on CPUs, GPUs, and Intel Xeon Phi
- Distributed by Sandia National Laboratories

<http://lammps.sandia.gov/>

---



LAMMPS rhodopsin benchmark (32,000 atoms).  
Courtesy Malakar et al. "Optimal scheduling of in situ analysis for large-scale scientific simulations." SC 2015.

# LAMMPS instrumentation with SENSEI

LAMMPS Input File

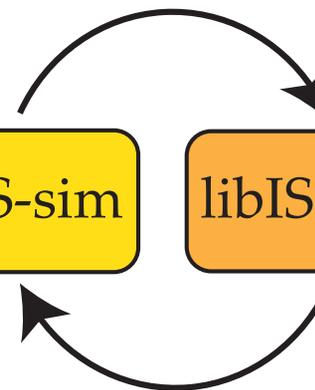
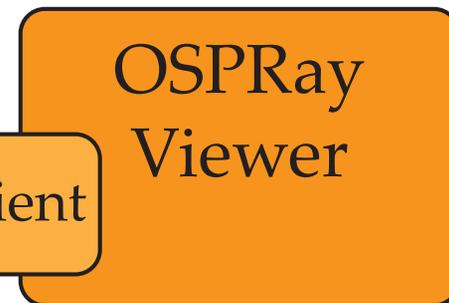
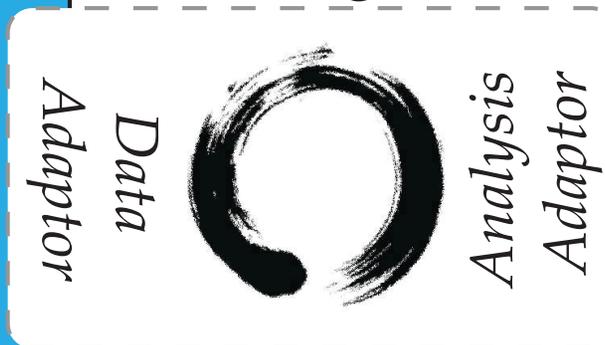
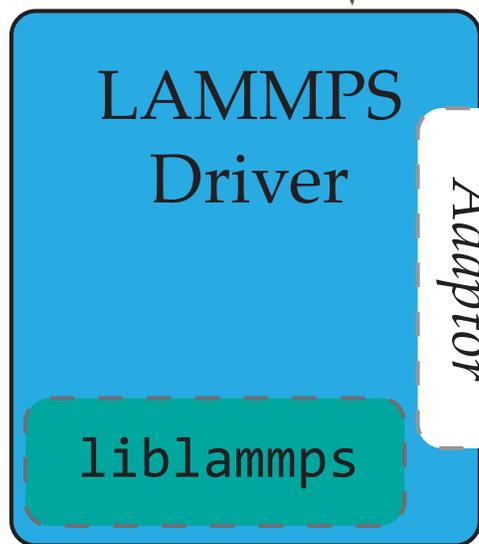


LAMMPS as a library

Bridge

In transit. Two concurrent jobs

Viewer decoupled from renderer



liblammmps

Adaptor

Data

Analysis  
Adaptor

libIS-sim

libIS-client

OSPRay  
Viewer

# Materials Science with LAMMPS

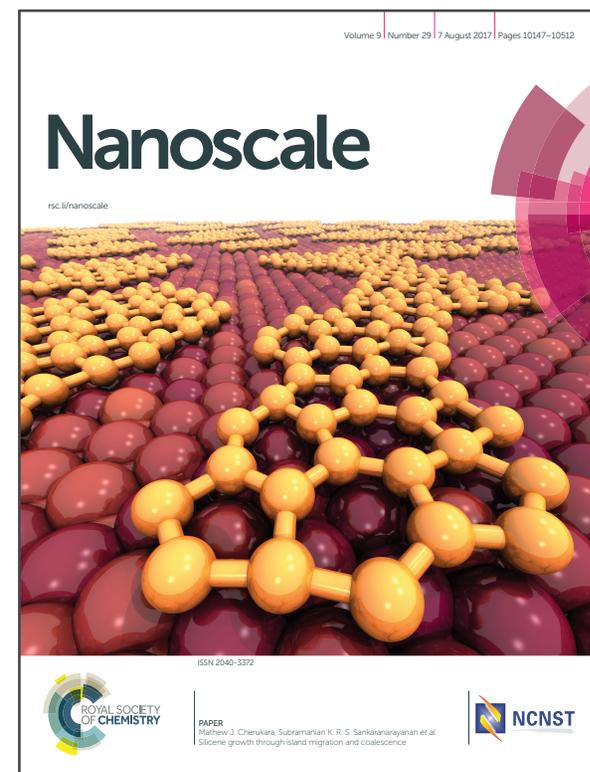
## Silicene: Mono-layer Silicon / Iridium Substrate

- Massively-parallel classical molecular dynamics (MD) simulations with LAMMPS
- Various temperature conditions
- Varying rates of silicene deposition
- Characterize material structure and growth

## Simulations were run on Mira at Argonne

162,000 iridium atoms

~6 Million total compute hours



Cherukara, Mathew J., Badri Narayanan, Henry Chan, and Subramanian Sankaranarayanan.  
"Silicene growth through island migration and coalescence." *Nanoscale* 9, no. 29 (2017)

Slide courtesy Joe Insley,  
Argonne National Laboratory

```
sci1952: sci1952 > work > bash
File Edit View Search Terminal Help
Pair      0.11205      0.12649      0.14685      3.9      82.74
Neigh     0            0            0            0.0      0.00
Comm      0.00049591   0.020846    0.035294    9.6      13.64
Output    0.00085711   0.00098503  0.0010428   0.0      0.64
Modify    0.0038671    0.0040505   0.004142    0.1      2.65
Other     0            0.0005049   0            0.33     0.33

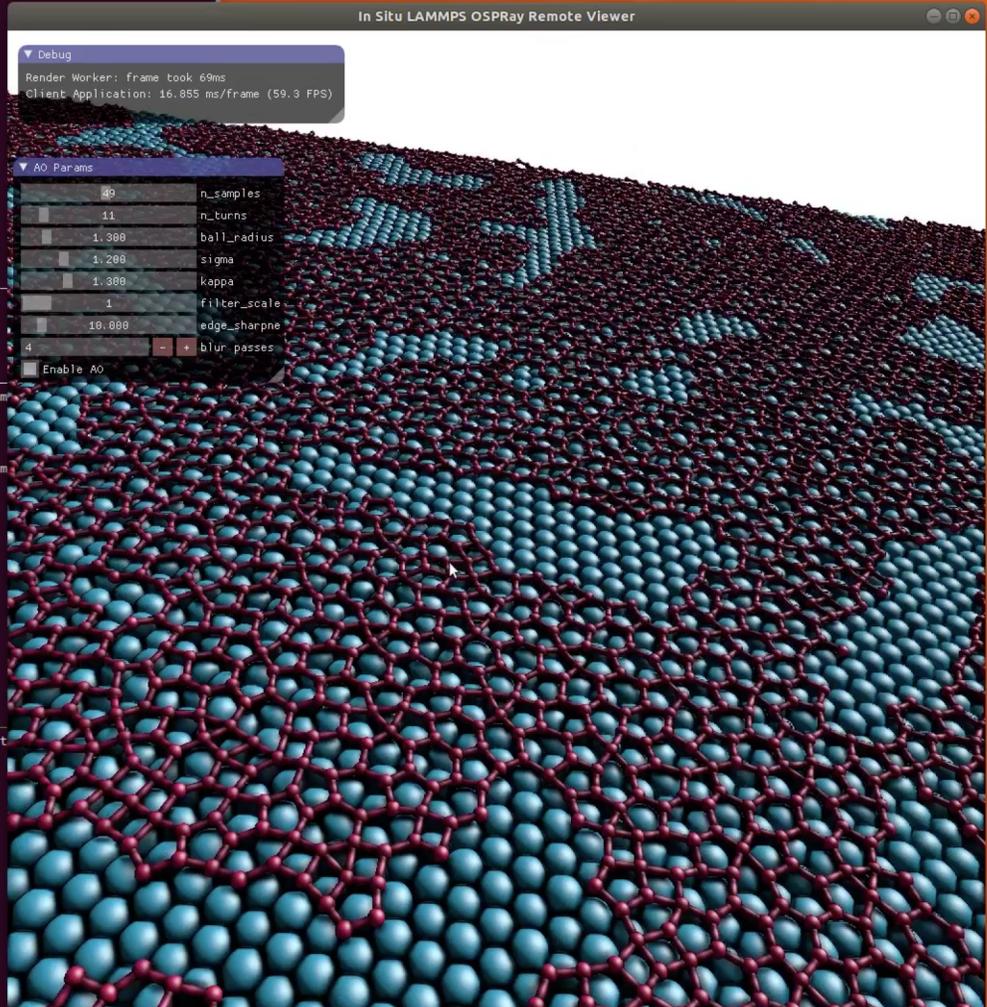
Nlocal:    28506.7 ave 28609 max 28388 min
Histogram: 1 0 1 1 0 1 0 0 0 2
Nghost:    16954.3 ave 17065 max 16816 min
Histogram: 1 0 1 0 1 0 0 1 1 1
Neighs:    556180 ave 568459 max 532930 min
Histogram: 1 0 0 0 0 1 1 1 0 2
FullNghs:  18490 ave 19059 max 17325 min
Histogram: 1 0 0 0 0 1 0 1 1 2

Total # of neighbors = 3337083
Ave neighs/atom = 19.5105
Neighbor list builds = 0
Dangerous builds = 0
WARNING: One or more dynamic groups may not be updated at correct point in timestep (../fix
WARNING: One or more atoms are time integrated more than once (../modify.cpp:275)
Setting up Verlet run ...
Unit style : metal
Current step : 57
Time step   : 0.0005

will@sci1952:~/repos/lammps_sensei_ospray/viewer/build$ I_MPI_FABRIC=shm I_MPI_SHM_LMT=shm m
is render worker -sim-host localhost -sim-port 29374 -port 6910 -bond 1 2.7
OSPRay with rank 0, world size: 1
Connecting over the network to the simulation
[0] DAPL startup: RLIMIT_MEMLOCK too small
Sending connect cmd, got MPI port name from open 'tag#0$description#sci1952$port#38231$ifnam
rank 0 running on sci1952
Now listening for client on sci1952:6910

will@sci1952:~/repos/lammps_sensei_ospray/viewer/build$ ./lammps_is_viewer -server localhost
Got world bounds = [(-0.0750253,-0.0592656,-5):(282.256,244.435,41.3664)]
[]

work1:~$ bash 2:~$ bash 3:~$ bash
will@sci1952
```





**BERKELEY LAB**  
Bringing Science Solutions to the World



U.S. DEPARTMENT OF  
**ENERGY**

# QUESTIONS ?

[sshudler@anl.gov](mailto:sshudler@anl.gov)

[srizzi@anl.gov](mailto:srizzi@anl.gov)

<https://sensei-insitu.org/>

